

Alaska SAR Facility Mass Storage, Current System

David Cuddy, Eugene Chu, and Tom Bicknell

MS 300-319
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109, USA
Phone: (818) 354-6277
Fax: (818) 393-5184
dcuddy@sakai.jpl.nasa.gov

Abstract

This paper examines the mass storage systems that are currently in place at the Alaska SAR Facility (ASF). The architecture of the facility will be presented including specifications of the mass storage media that are currently used and the performances that we have realized from the various media. The distribution formats and media will also be discussed. Because the facility is expected to service future sensors, the new requirements and possible solutions to these requirements will also be discussed.

Introduction

Synthetic Aperture Radar (SAR) is an imaging radar technique that achieves high resolution through synthesizing the performance of a large aperture radar with a small aperture by processing the data from multiple samples of the beam footprint across the target in the azimuth direction. Typically the SAR data is collected from spacecraft or aircraft. To achieve high resolution and wide coverage, large volumes of data are collected from the SAR, and large data sets are produced when the data is processed into images.

The ASF is a National Aeronautics and Space Administration (NASA) sponsored project to collect data from space borne SAR instruments, to process this data into SAR images, and to operate an active archive in support of scientific investigations which use these images for the study of geophysical processes. Currently, the ASF is collecting SAR data from two satellites: the European Space Agency's (ESA) first Remote-Sensing Satellite (ERS-1) and the Japanese space agency's (NASDA) first Earth Resources Satellite (JERS-1). ERS-1 was launched in July of 1991, and JERS-1 was launched in February of 1992. Both satellites have been transmitting SAR data steadily to the ASF since their initial validation period. The ASF will evolve to support ERS-2 which is ESA's second in the series of Remote-Sensing Satellites with a launch in late 1994 or early 1995, and to support the Canadian Space Agency's (CSA) RADARSAT which is scheduled to launch in early 1995. The ASF was designed and built by Jet Propulsion Laboratory and is located at the University of Alaska in Fairbanks (UAF) and is operated by the UAF. The ASF is now one of the Distributed Active Archive Center (DAAC) sites by the Earth Science Data and Information System (ESDIS) project.

ASF Architecture

The overall functional diagram of the ASF is shown in Figure 1. The ASF consists of four major components [1]: the Receiving Ground Station (RGS), the SAR Processor System (SPS), the Geophysical Processor System (GPS), and the Archive and Operations System (AOS). The RGS functions are to track the satellites with a 10 meter dish, to receive and record the SAR data sent by the satellites on high density recorders for both the local archives and for the

respective flight agencies. The SPS functions are to read and decode the data that has been recorded on high density tapes by the RGS, to process the SAR signal data into SAR image data products, and to deliver these products to the archives. The GPS functions are to create geophysical products such as ice motion vectors, ice classification images, and wave products from the SAR images, and to deliver these products with their metadata to the archives. The AOS functions are to manage the archive, to provide a user interface for searching the catalog of the metadata and for ordering data from the archives, to manage the data distribution, and to provide mission planning capabilities including the ability for a user to request a data acquisition.

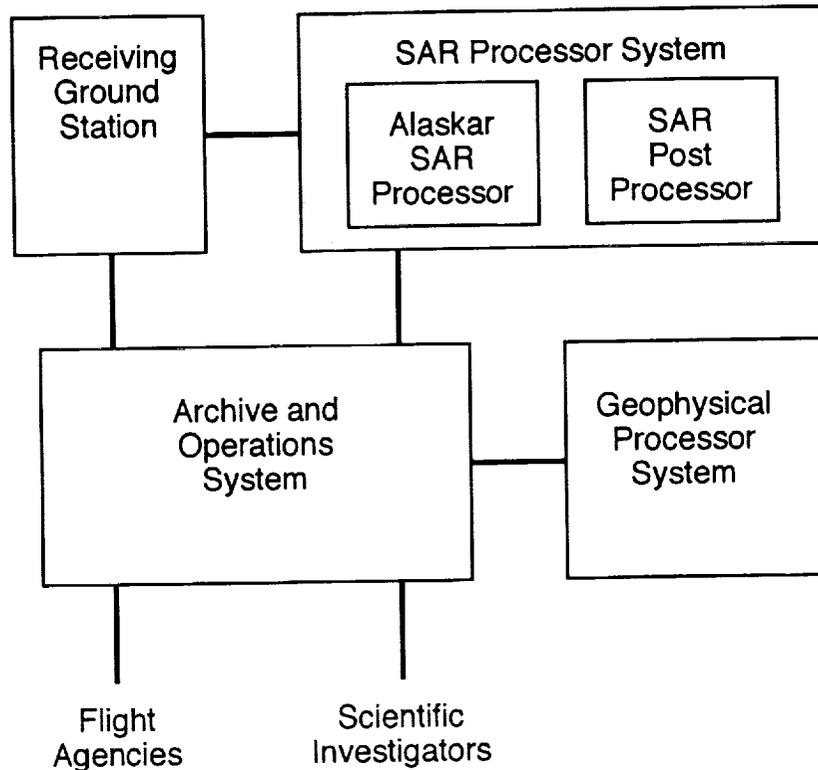


Figure 1: The Alaska SAR Facility Functional Block Diagram.

Radar Instruments

The ERS-1 satellite generates data at 105 Megabits per second during a data take which can last up to 15 minutes. The ASF receives multiple data takes per day. For 60 minutes of collected SAR data, the volume amounts to 47.3 Gigabytes (105x60x60 Megabits) of the raw data. The JERS-1 satellite has a lower data rate of 65 Megabits per second, but it also has an on-board recorder which can concurrently dump data recorded from anywhere around the earth with the data collected in real time.

For standard processing of ERS-1 data, each minute of data generates about 4.5 images, or 9 images from each 2 minutes of data. Each image is 8kx8k pixels covering an area on the earth of 102.4km x 102.4km with pixel spacing of 12.5 meters in each dimension. One of the standard products derived from each image is a 1kx1k low resolution image generated by performing an 8x8 average of the full resolution image. For each image there are also metadata that describe the image and the processing that produced the image. Each metadata set is about 30 kilobytes (KB) of information. The total storage required for each image is about 68.2 Megabytes (MB), or 20.5 Gigabytes (GB) per 60 minutes of data. For a 1000 day mission, this translates to 36.7 Terabytes of input data and 20.5 terabytes of output data. Mass storage is one

of the major cost drivers in the operation of the ASF, as it will be for all SAR data systems. This has become especially true with computing systems rapidly increasing in power. A few years ago, the cost of high volume and high speed storage systems were secondary to the cost of high performance computing systems. In recent years, the relation has reversed.

The RGS

The RGS records the telemetry and SAR signal data to many combinations of high density recorders including: two Honeywell HD-96 for data sent to Japan, one Thorn-EMI Digital tape recorder for data sent to ESA and six AMPEX DCRSi recorders for the local archives and for inputs and outputs of the subsystems that operate on the data. Typically at transmission time, the RGS will record two DCRSi tapes, one for the archives and one for the working copy, and one tape for the respective flight agency. Because the JERS-1 has an on-board recorder, it can transmit dual data streams - one from the on-board recorder and the other from the SAR instrument in real-time, so the RGS will record data to both Honeywell recorders simultaneously as well as to two DCRSi drives.

The SPS

The SPS consists of two subsystems: the Alaska SAR Processor (ASP) and the Post Processor (SPP). The ASP is composed of a custom hardware for reading, decoding, and processing of SAR data from DCRSi tapes and a MASSCOMP computer which controls the custom hardware and the DCRSi tape drives. It uses one DCRSi drive to input SAR data and one to record the output results. The ASP reads ancillary data from the DCRSi to generate processing parameters which are used to set up the custom hardware for processing the SAR data.

The SPP is composed of two computers that work together for specialized input and output and for making the products compatible with VMS-RMS file system. A DEC MicroVAX is used to control all of its functions, while an APTEC IOC-24 is used to perform some of the high data rate functions. Its functions include reproducing high volume data from the DCRSi, receiving averaged image data from the ASP, converting all data into the VMS format for the AOS, and recording image data onto films. The SPP uses a disk system shared with the AOS to quickly exchange large volumes of data. In addition, the SPP uses local storage to keep its copy of metadata and ancillary information.

The GPS

The GPS, which produces the higher level geophysical products, is implemented on a SUN-4 workstation with a high-speed array processor. It communicates with the ACS via an Ethernet connection. It needs enough local storage for its own database of metadata and for work space to handle input image data and output geophysical products. This space amounts to 3 GB of magnetic disk storage.

The AOS

The AOS consists of two subsystems: the Archive and Catalog Subsystem (ACS) and the Mission Planning Subsystem (MPS). The ACS maintains the archive and the catalog of the archive, and the ACS provides the user interface so that they can search the catalog, order data from the archives, and request data be acquired from a satellite. ACS is providing the IMS and DADS functions in the ESDIS DAAC model. The MPS provides the mission planner with tools to perform conflict resolution of satellite schedules and to create weekly operations schedule of data acquisition times.

The MPS is implemented on a VAXstation 4000/90 and the ACS is implemented on a VAX 8530 which will soon be augmented with another VAXstation 4000/90 to perform its database operations. The ACS has a large volume (24 Gigabytes) of magnetic disk storage to serve several storage requirements. The largest use is to cache data as it is moved to and from the optical disk jukebox, which stores all of the low resolution data and all of the higher level products in the near-line archive. Other uses for the magnetic disk include staging area for data that is to be transferred to the user electronically, housing of the database for the archive catalog, and providing work space for some of the auxiliary processing. For digital image distribution, the ACS has two 9-track tape drives with multiple density capabilities (800 bpi, 1600 bpi, and 6250 bpi) and two 8-mm tape drives. The ACS can make prints of low resolution images on a Lasertechnics printer. One of the major auxiliary processing that ACS must perform is to transform the images onto a map projection such as Universal Transverse Mercator or Polarstereo projections by using an array processor and a large portion of its 96 Megabytes of memory.

ASF High Density Storage Systems

The ASF currently uses two types of high density tape recording formats for data storage in the RGS and the SPS. One format used in the Honeywell and EMI high density data recorders (HDDRs) is the familiar reel to reel 1 inch wide high density data tape (HDDT) with multiple linear data tracks that are recorded by moving the tape over a fixed head assembly. The other format is the AMPEX DCRSi which uses a rotating head assembly that records data tracks across the width of the tape as it is moved over the heads.

The reel to reel HDDTs were traditionally used for recording high rate data collected from instrumentation and satellite systems. The tapes are one inch wide, and each reel contains approximately 9600 feet, providing approximately 15 minutes of recording time at the maximum data rate. The reels are approximately 14 inches diameter and each reel of tape weighs about 20 pounds. Both the EMI and Honeywell drives are capable of recording data at various rates from a minimum of less than 1 megabit per second to a maximum of about 150 megabits per second on up to 32 data tracks. The drives can record data onto each track with a maximum density of 33 kilobits per inch. With all 32 tracks in use, this provides an aerial density of approximately 1 megabits per square inch. This translates to approximately 12 GB of data capacity per reel. In practice, some of the data tracks are used for recording error checking and correcting codes (ECC), so the effective data density and capacity are somewhat less.

The recorders input data via a serial data line and use a separate clock line that synchronizes the data. The drive electronics track the input clock and set the tape speed and density accordingly to keep up with the data. The recording rates are changed by varying the speed of the tape motion; the drives are capable of moving the tapes at a minimum of 15/16 inch per second (ips) to a maximum of 120 ips, in binary increments. As the tapes require different recording and playback equalizations at each different speed, a different set of equalization and biasing circuits in the recorder are required for each speed of operation. As a new tape speed is selected, a different set of recording electronics, tuned for that speed, is selected. Each set of recording electronics can tolerate small variations of data rates, and the recorder can compensate for some of these variations by slewing the tape speed and changing the data density on the tape. However, large variations of data rates will cause the recorder to select the next higher or lower tape speed, and the associated recording electronics. But because of the size and weight of the reels, the recorders can not change the speed of the tape motion very quickly. Therefore, during recording, the input data rates must remain fairly constant since some data will be lost as the recorder selects a new speed and a different set of recording electronics.

When preparing for a recording, the tape drives must be started at least 15 seconds before the input data becomes available in order to spin up the reels to move the tape at the proper speed. In practice, however, much more lead time is used for the spin-up and pre-roll functions in

order to insure that no data is lost. Similarly, during playback, the tape motion must be started from a point well before where the data begins to allow the reels and tape to reach their correct speeds. Then the data is output by the recorder at a constant rate with a clock signal, and the device receiving the data must be capable of ingesting the data at that rate. Typically, in order to process the data, the receiving device is a general purpose computer system. These can only accept the data at rates far lower than when it was recorded. In addition, the computer most likely could not be performing any other tasks at the same time or it could be diverted from the data receiving task long enough to lose some of the data. Often in practice, when a tape is played back at a much slower speed than when it was recorded, the playback process is not reliable, and is subject to very high error rates.

In order to facilitate location of data on a HDDT, time codes are recorded onto the tape with the data. During playback, a time code decoder is used to find the approximate starting location of the data area of interest. Then, as the recording is played back to the processing computer, it will contain the desired data as well as unneeded information. It is the task of the processing computer to separate these by locating the desired data.

The ASF provides data recording services for ESA and NASDA when their respective satellites are within its station masks using the tape formats that each agency uses. The JERS-1 satellite uses an on-board recorder to store some of the data that it collects. The recorder does not rewind its tape when it dumps the recorded data, so it is dumped in reversed order from when it was collected. The two Honeywell recorders have the ability to play its data in reverse order as well, so the reversed data that was recorded onto it can be played back in the correct order.

The ASF uses the AMPEX DCRSi cassettes for storing its own long term archives of ERS-1 and JERS-1 satellite data and all processed data generated by the SPS by processing the SAR data. The digital cassette format is a relatively new entry in the field of high density data recording. The DCRSi is derived from an AMPEX broadcast video cassette format. The tape itself is 1 inch wide and each cassette contains about 1600 feet, providing approximately 1 hour of recording time at its maximum data rate of 107 megabits per second. This translates to approximately 48 gigabytes of data capacity. The record/playback head assembly consists of 6 heads mounted on a drum which rotates in a direction perpendicular to the direction of tape motion. As the tape is moved across the head assembly, successive "swipes" of data are recorded across the width of the tape. Each block receives its own address during recording, so it can be referenced during playback. The recorder also records a linear track onto the tape containing coarse address markers which aid in location searches.

The DCRSi moves the tape at only one speed, the maximum speed, and records data onto the tape at one density, the maximum density. The recording heads scan across the width of the tape, recording 4356 bytes, or 34848 bits, of user data in each block, with a linear density of 606 blocks per inch. This represents an aerial density of over 21 megabits per square inch. This is effective data density, as actual density is somewhat higher to record ECC data, time codes, and block addresses in addition to the user data.

The DCRSi recorder uses a front-end buffer to catch input data and spool output data. During the recording process, as the buffer becomes half filled with the input data, the tape drive pre-rolls the tape and begins recording the buffer data at its maximum speed. As the buffer empties, the tape motion stops, and the tape is repositioned, ready to start moving again when the buffer becomes half full. The tape motion is simple and quick enough that it can be accomplished before the buffer fills up. During playback, the drive first plays data into the buffer. If the destination device cannot accept the data fast enough, the buffer will fill up, and the tape will be stopped and repositioned. As the buffer reaches half empty, the tape is pre-rolled and started again, filling the buffer. This allows the external data rate to vary continuously between 0 and the maximum rate during recording or playback with no restrictions on how often the rate varies. In addition, there is no need to start recording on the tape before the input data becomes available, as there is with the HDDTs. The recorder can respond to incoming data and outgoing data requests with almost no latency. This allows the data to be recorded at, and later to be

retrieved from, exactly the specified location on the tape. This also has the benefit of helping to make more efficient use of the tape.

The flexible data flow characteristic of the DCRSi is established once the recorder is put into the proper recording or reproducing mode. This step requires that the recorder move the tape to the requested block, which could require up to two minutes, depending on what position the tape was in initially. But once the initial seek is completed, the recorder response to incoming or outgoing data is nearly instantaneous. Effectively, the DCRSi behaves like a 48 gigabyte disk drive with long seek times, fast transfer rates, but no latency times.

The versatility of the DCRSi makes it much easier to use than the reel to reel HDDRs. Its capabilities have made it the only recorders used in the operational data processing in the SPS. The DCRSi cassette itself is capable of storing about four times as much data as a full reel of HDDT, but it requires approximately one-fourth the physical volume for storage. This makes it a much better medium for long term archives in the ASF. The one function that the DCRSi cannot perform for the ASF is to play its data in reverse, as the linear track recorders can. So in order to process the JERS-1 data collected from a dump from the on-board recorders, the data must first be dubbed onto a DCRSi cassette from a Honeywell HDDR playing in reverse.

The DCRSi does share some inconvenient features with the older HDDRs. Because of the high data rates of these types of recorders, they cannot be easily interfaced to a general purpose computer for access of their data. The DCRSi currently requires a customized interface for communicating with any other device. On the RGS, the standard serial data and clock lines adequately communicates with the drives. On the ASP, a custom interface was designed in-house to make use of the DCRSi's parallel interface. On the SPP, a special interface was designed by AMPEX to communicate with the APTEC IOC-24, also using the parallel interface.

The ASF is currently investigating a high speed SCSI interface for the DCRSi. It consists of a DCRSi interface installed into the VME chassis of a small computer system based on a SUN SPARC processor. The system also includes a fast SCSI interface for interfacing to general purpose computers. This will enable the DCRSi to be attached to most computer systems with a fast SCSI interface.

The ASF currently operates 6 DCRSi drives. The RGS requires at minimum 1 DCRSi drive, although the normal operation is to record on two drives at down link time to concurrently make both the archive and working copy. The SPS requires three DCRSi drives, one for input and one for output during SAR image processing and one for playback of the processed data. The sixth drive is kept for a working spare which allows quick change when a drive fails. A parallel switch helps to easily configure which drive is connected to which port of a specific computer.

To store all of the data on-line would be prohibitive and the archive strategy to this date is to store all large volume data sets on high density recorders in the off-line manner. The location (tape identity and addresses on the tape) and other key information are kept in an on-line data base in the ACS.

The ACS uses a jukebox with a capacity of 89 platters of Write-Once-Read-Many (WORM) laser disks for storage of low resolution data. With a 2 GB capacity per platter, the juke box has a capacity of 178 GB. The jukebox contains two drives each capable of addressing only one side of the platter at a time. The robotic mechanism in the jukebox moves the platters between their storage bins and the drives, and positions the selected side of the platter in the drive. To prevent excessive robotic action, the input data is cached on magnetic disks on the VAX, and periodically flushed to the jukebox. Since the disk platters are normally stored in their bins, the initial access time to each platter is approximately 30 seconds, the time required by the robotic mechanism to retrieve the platter, insert it into the drive, and spin it up to speed. Then, the optical disk drive behaves like a slow magnetic disk drive; they are capable of transferring data at 250 kilobytes per second, with seek times on the order of tens of milliseconds.

The ACS has a total of 24 GB of magnetic storage to provide storage space for caching of data to the archive, to provide storage space for the catalog and the needed work space for the database system, to provide working space for geocoding of images, to provide cache space for staging of data to users and to the GPS, and to provide work space to all users. It uses 6 GB of magnetic disks for caching data for the optical disk jukebox. In addition, it shares 8 GB of disk storage with the SPP for buffering of data exchanged in between the two systems.

Initially, the ACS used physically large disk drives connected to the VAX through VAXBI controllers using a proprietary DEC interface. These drives and controllers were costly, not very reliable, and consumed large amounts of physical space and power. The ACS has now been upgraded to new disk drive technologies. In the last two years, disk manufacturers have released a number of technologically advanced disk drives into the market. These drives are physically smaller, transfer data at higher rates than older drives, and are far less expensive in terms of cost per unit of storage. These drives also use a new version of the Small Computer System Interconnect (SCSI-2) interface standard for connection to the host computers. This open standard interface allows any of these drives to be connected to any host computer with a SCSI adapter.

The SPP uses SCSI-2 disk drives almost exclusively for its operations. Its 8 GB of shared disks are composed of 4 SCSI drives which have dual host attachments. The SPP writes the data it produces onto the shared disks for the ACS to retrieve, and the ACS writes processing requests and data onto the disks for the SPP to read. The two systems are using the disks in a static dual-port manner; one disk is configured to allow only the ACS to write to and the SPP to read from while the rest are configured to allow the SPP to write to and the ACS to read from. The VAX clustering feature on the VMS operating systems on both the MicroVAX and the 8530 can be enabled to perform full active dual-porting of the shared disks. However, it was found that the overhead of VAXclustering was too great, and that the current configuration is more than adequate. The SPP also keeps its own repository of all full resolution metadata on its local storage of 4 GB of SCSI disk drives.

The ASP MASSCOMP currently uses disk drives with the old Storage Module Device (SMD) interface. There are a total of 1 GB of storage used for operations. The ASP is in the process of replacing the MASSCOMP with one of the newer workstation based systems which will be much faster, easier to maintain, and will also use the more efficient SCSI disk drives for storage.

Data Distribution

The ASF is responsible for distribution of data on various media to the scientific users. Currently, the ASF delivers data on the following media: nine track magnetic tape at 1600 bytes per inch (bpi) and 6250 bpi, 8 mm tape at 2.3 gigabytes or 5.0 gigabytes per tape, file transfer of all files, and hard copies of images on low or high resolution films. The data is stored and distributed in a format that is in compliance with the specifications of the Committee on Earth Observations Satellites (CEOS) [2], to which ESA, NASDA, NASA, and CSA have agreed as the format of data exchange. The data that is put on tape also have a CEOS volume description that identifies the contents of the tape(s). The CEOS specifications allow for leader and trailer files which describe the details of the data, the processing, the sensor, and the satellite. The hardcopy and film products also are in compliance with the guidelines set by the CEOS specifications.

The digital data distributed on the various recording media include low and high resolution multi-look detected image data, high resolution single-look complex data, and reformatted SAR signal data. Initially, the ASF supported only the 9 track magnetic tape as its sole form of data distribution media. These tapes are usually packed with 3600 feet per reel, providing up to 270 MB of storage in 6250 bpi and up to 69 MB at 1600 bpi. In the higher density, this provides the ability to record about 80 low resolution detected images, two high resolution detected

images, or one complex image. The reformatted SAR signal data files need to be spread across two reels, or volumes of tapes.

More recently, the ASF began distributing data on the 2.3 and 5 GB 8 mm cassette tapes. Because of its high capacity, a single 8 mm cassette can usually store all the data requested by any user. It requires less operator assistance since it is rare to need to change tapes to make more than one volume for any request. The media is also much smaller and lighter than the 9 track tape, making it easier to handle store, and ship. In addition, all of the 8 mm tape drives available are produced by the Exabyte Corporation, which has created a de-facto standard in the 8 mm format. All drives also use the SCSI standard interface, making it readily usable by most computing platforms.

The 8 mm tape format does have some drawbacks compared to the old 9 track tapes. First, its drive mechanism is much slower than those of the 9 track machines. Typically, any given tape motion on the 8 mm drive can take up to an order of magnitude longer to complete than the 9 track drives. Then, the data transfer rates of the 8 mm drives are about half as fast as the 9 track drives. Finally, the one operation that is the most time consuming on the 8 mm drives is the creation of a new tape file. The Exabyte tape format uses a very large header for each file. Each header can be several MB in length, and requires on the order of tens of seconds to create. Using the example of a CEOS formatted tape, each product contains a leader file, the data file, and a trailer file, each of which would require a physical file header to be created by the drive. In addition, the CEOS tape also has a volume directory file, and a null descriptor file. So the overhead of creating a CEOS formatted 8 mm tape usually requires more time than transferring the actual data. In contrast, file headers are very efficient on the 9 track drives, and are usually executed very quickly.

The hard copies of image data are distributed on two types of films. The low resolution data are recorded onto dry silver films using a LASERTECHNICS 300D printer. The high resolution data are recorded onto Kodak photographic films using a ColorFIRE 240 film recorder made by MacDonald Dettwiler and Associates (MDA), now Cymbolic Sciences Incorporated (CSI).

The 300D printer takes 8 bit image pixels as its input. The data is used to modulate a Bragg crystal which in turn modulates the intensity of a laser beam that exposes the pixels on the dry silver film. The modulated beam is scanned across the width of the film by an oscillating mirror. The film is developed by heat rollers within the printer, providing a nearly instant hardcopy on the image. The printer is capable of writing up to 2048 pixels across a 8.5 inch wide film, and up to 1500 lines of pixels. The transparency film provides 128 effective gray shades, or 7 bits of dynamic resolution. The paper film provides 64 shades, or 6 bits of resolution.

The ColorFIRE 240 is capable of making use of full 24 bit color (8 bits red, 8 bits green, 8 bits blue) image data to generate color films. However, the SAR image data that the ASF currently produces has only one channel at 8 bits per pixel, so only black and white films are being generated from existing data. The recorder is capable of writing up to 8800 pixels across the 240 mm (9.44 inch) film, and up to 9600 lines of pixels. It also uses a Bragg crystal to modulate the intensity of a beam of light which exposes the film. The beam is directed across the width of the film by a rotating mirror riding on air bearings. This allows the printed image to have extremely accurate geometric fidelity. The ColorFIRE 240 in the ASF is set up to also provide very accurate radiometric accuracy while maintaining the full 8 bits of dynamic resolution.

RADARSAT Support

The ASF facility was originally sized to collect and process 5 minutes per day of ERS-1 data, 10 minutes per day of JERS-1 data, and 30 minutes per day of RADARSAT data. During the first year of ERS-1 operations, ASF collected on the average more than 30 minutes of data per day and for JERS-1 the volume has also exceed the original sizing. For RADARSAT, the system is now being specified with a requirement of 80 minutes of processed data from 120 minutes of

collected data. This translates roughly to 5 terabytes of processed and 33 terabytes of raw data per year. The remaining unprocessed data is to be collected for other stations who will be responsible for processing their own data.

Currently the ASP processor at ASF is being upgraded to have the ability to process 60 minutes of data per day. For the RADARSAT, additional processors will be delivered to ASF so that the processing throughput will be increase significantly to keep up with the data volume. The storage capacity is quickly being consumed even at the processing rate of 40 minutes per day. The storage of on-line data will exceed capacity the current archive system within 6 months, and ASF is investigating different technologies and archive strategies which can alleviate the expected overflow in the near future, but the studies are looking also at the big picture of the many years of data to come and also at the ESDIS schemes to handle storage with their DAAC contract.

Conclusion

SAR data is a very high volume form of data which requires a processing and archive facility which can accommodate multiple terabytes of input and output data per year. The data distribution also must be able to handle large volumes of data to a very diverse community of users whose requirements on the data are almost unique on a user by user basis. The ASF has been able to meet the demands for the large volumes of data by employing not only a variety of storage media but an archive strategy that tries to keep the more frequently accessed and smaller data sets either on-line or near-line while the less frequently accessed and larger data set are kept in the off-line, high volume storage media. Granted, the system is tasked rather heavily, but it has exceeded its original requirements and will be able to grow and evolve with the ever increasing requirements of future.

Acknowledgments

This work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] R. Berwin, D. Cuddy, J. Hilland, and B. Holt, "Design, Test, and Applications of the Alaska SAR Facility," Space Technology, Vol. 12, PP 91-104, (1992).
- [2] CEOS WGD on SAR Data Standards, "SAR Computer Compatible Tape Format Specifications", Rev 2.1, January 1, 1992.

